Background of the Invention

The present invention relates to local area networks (LANs) which enable devices with computational ability to communicate with each other and, in particular, to a wireless LAN in which the devices communicate by means of radio transmissions.

In recent years the personal computer has become an increasingly important tool in business and commerce and many workers now spend a good portion of their working day operating such computers. Similarly, business organisations are increasingly structuring their businesses to not only enable, but to oblige, their workers to access information by means of a personal computer or equivalent terminal, which is connected to a local area network which extends around or through the office environment.

Hitherto such local area networks have been provided either by electrical conductor or optical fibre and this requires the office premises to be extensively cabled. This cabling must be adjusted if, for example, partitions within an office are to be adjusted. In addition, the cabling required for a classroom or tutorial arrangement where a large number of personal computers are intended to be operated within a small areas, can be quite substantial.

Furthermore, an increasing trend in recent times has been the sale of mobile or portable devices with computational ability. These include both laptop/notebook and handheld computers. Whilst the primary impetus for the purchase of such a computer is the ability to use its computational power outside of the normal office environment, once a portable computer has been purchased, the desire arises to use the portability within the office premises so as to allow the user of the portable computer to take the computer with him and use it in the closely adjacent offices of colleagues, for example, and yet still be able to access the LAN of the business organisation, which may be spread over several adjacent buildings in "campus" style.

While this is possible by means of plug-in connectors which enable the portable computer of one operator to be plugged into the office LAN at any particular location, it is generally inconvenient since the LAN may not provide for two or more points of connection within a single office, the portable computer loses its portability, and so on.

Accordingly, the need arises for a LAN to which such portable

devices can be connected by means of a wireless or radio link.

Such wireless LANs are known, however, hitherto they have been substantially restricted to low data transmission rates. In order to achieve widespread commercial acceptability, it is necessary to have a relatively high transmission rate and therefore transmit on a relatively high frequency, of the order of 1 GHz or higher. As will be explained hereafter, radio transmission at such high frequencies encounters a collection of unique problems.

One wireless LAN which is commercially available is that sold by Motorola under the trade name ALTAIR. This system operates at approximately 18 GHz, however, the maximum data transmission rate is limited to approximately 3-6 Mbit/s. A useful review of this system and the problems of wireless reception at these frequencies and in "office" environments is contained in "Radio Propagation and Anti-multipath Techniques in the WIN Environment", James E. Mitzlaff IEEE Network Magazine November 1991 pp. 21-26.

This engineering designer concludes that the inadequate performance, and the large size, expense and power consumption of the hardware needed to adaptively equalize even a 10 Mbit/s data signal are such that the problems of multipath propagation cannot thereby be overcome in Wireless In-Building Network (WIN) systems. Similarly, spread spectrum techniques which might also be used to combat multipath problems consume too much bandwidth (300 MHz for 10 Mbits/s) to be effective. A data rate of 100 Mbit/s utilizing this technology would therefore consume 3 GHz of bandwidth.

Instead, the solution adopted by Motorola and Mitzlaff is a directional antenna system with 6 beams for each antenna resulting in 36 possible transmission paths to be periodically checked by the system processor in order to locate the "best quality" path and "switch" the antennae accordingly. This procedure adds substantial bulk and cost to the system. This procedure is essentially the conversion of a multipath transmission problem into a single path transmission environment by the use of directional antennae. Objects and Summary of the invention

The object of the present invention is to provide a wireless LAN confined in a multipath transmission environment having a high bit rate, such even through that the reciprocal of the data or information bit rate (the data "period") is short relative to the time delay differences between

significant transmission paths.

According to one aspect of the present invention there is disclosed a wireless LAN comprising a plurality of hub transceivers each connected together to constitute a data source and/or destination, and a plurality of mobile transceivers each able to communicate by radio transmissions with any one(s) of said hub transceivers within a predetermined range, wherein each of said mobile transceivers are connectable to, and able to be powered by, a corresponding portable electronic device with computational ability, said radio transmissions have a frequency in excess of 10 GHz, and all the transceivers are configured to receive and transmit in a multipath transmission environment, the reciprocal of the information bit rate of said transceiver's transmission being short relative to the time delay differences between significant ones of the transmission paths of said multipath transmission environment.

According to another aspect of the present invention there is disclosed a peer-to-peer wireless LAN having a plurality of mobile transceivers each able to communicate by radio transmission with any other like transceiver within a predetermined range, wherein each of said mobile transceivers are connectable to, and able to be powered by, a corresponding portable electronic device with computational ability, said radio transmissions have a frequency in excess of 10 GHz, and all the transceivers are configured to receive and transmit in a multipath transmission environment, the reciprocal of the information bit rate of said transceiver's transmission being short relative to the time delay differences between significant ones of the transmission paths of said multipath transmission environment.

According to a still further aspect of the present invention there is disclosed a method of transmitting data between at least one hub transceiver and a plurality of mobile transceivers within a predetermined cell range or between said mobile transceivers, wherein said data transmission is a multipath transmission having a frequency in excess of 10GHz, each said mobile transceiver is connected to, and is powered by, a corresponding portable electronic device with computational ability, and the reciprocal of the information bit rate of said transmissions is short relative to the time delay differences between significant ones of the transmission paths of said multipath

transmission environment.

Preferably, transmission is enhanced by the use of one or more of the following techniques, namely interactive channel sounding, forward error correction with redundancy sufficient for non-interactive correction, modulation with redundancy sufficient for interactive error correction by re-transmission of at least selected data, and the choice of allocation of data between sub-channels.

The radio transmission is also preferably divided into small packets of data each of which is transmitted over a time period in which the transmission characteristics over the predetermined range are relatively constant.

The encoding of the data is preferably carried out on an ensemble of carriers each costituting a sub-channel and having a different frequency with the modulation of each individual carrier preferably being multi-level modulation of carrier amplitude and/or phase (mQAM). The modulation family mQAM includes amplitude shift keying (ASK), multi-level ASK (mASK), permutation modulation, binary phase shift keying (BPSK), multi-level phase shift keying (mPSK), amplitude phase keying (APK), multi-level APK (mAPK), and the like.

Brief Description of The Drawings

Embodiments of the present invention will now be described with reference to the drawings in which:

Fig. 1 is a schematic plan view of an office illustrating multipath transmissions of radio frequencies of at least 10 GHz caused by reflections;

Fig. 2 is a graph of received power as a function of time, for an impulse transmission, illustrating the received signals of reduced magnitude which are delayed owing to the possibility of multiple path transmission;

Fig. 3 is a graph of the received amplitude of steady state signals as a function of the transmitted frequency, this characteristic itself being time dependent;

Fig. 4 is a schematic diagram illustrating a local area network including a plurality of hubs each of which is able to communicate with mobile transceiver(s) within a corresponding cell;

Fig. 5 is a schematic block diagram of the circuit arrangements within each hub and mobile transceiver;

Fig. 6 is a more detailed block diagram illustrating part of the

mobile transceiver of Fig. 5.

Fig. 7 is a more detailed block diagram of the framing, FEC and modulator section 32 of the transmit path of the mobile transceiver of Fig. 6;

Fig. 8 is more detailed block diagram of the framing, FEC and demodulator section 32 of the mobile transceiver of Fig. 6; and

Fig. 9 is a more detailed block diagram of the mm-wave transmitter 36 and receiver 35 of the mobile transceiver of Fig. 6.

In schematic form, Fig. 1 illustrates a room 1 in a typical office environment which includes items of furniture 2 and a transmitter 3 and receiver 4. For radio transmissions at a frequency in excess of 10 GHz, a multipath mode of transmission from the transmitter 3 to the receiver 4 occurs. Reflections from the walls (and floor and ceiling) of the room 1, items of furniture 2, and the like, within the room 1 cause the multiple path transmissions.

As illustrated in Fig. 2, the effect of the multiple path transmissions is that the receiver 4 receives an undelayed signal 5 which has travelled directly from the transmitter 3 to the receiver 4, and a number of delayed signals 6 which are received at a time after receipt of the undelayed signal 5. The magnitude of the delayed signals 6 is usually somewhat attenuated. Under some conditions, the magnitude of the undelayed signal 5 can be attenuated also, sometimes by more than some delayed signals 6.

As a consequence of the delayed signals 6, it is necessary for the length of time during which a single symbol is transmitted (the symbol period) to be substantially longer than the delay time in order that the received echoes of a first symbol not mask the receipt of a subsequent symbol. This requirement has provided a severe upper limit to the rate at which data can be transmitted in such an environment.

Furthermore, as illustrated in Fig. 3, the office environment is by no means a good one for radio transmission. Fig. 3 illustrates a typical channel characteristic over a short time period illustrating the magnitude of the received signal as a function of frequency in the 1 GHz band between 60 and 61 GHz. It will be seen that the received amplitude is by no means constant and, in particular, at various frequencies fading occurs. Furthermore, as indicated by dotted lines in Fig. 3, the frequency at which fading occurs varies as a function of

time because of movements within the room. Such a communication channel is called a time varying frequency selective fading channel.

Similar, but different, communications channels are known in both telephone and long distance radio communications and various strategems, generally known as equalisation, are used to overcome the problems such channels present. However, in these fields since such fading is due to changes in temperature, or atmospheric conditions, once such telephone or long distance radio communication channels are established, the fading characteristic changes relatively slowly. Also in telephone applications advantage of the fact that channel degradation is generally low near the centre of the channel, can be taken when arranging the equalisation. This is not the case in an office or indoor environment.

Rather, in the above described office environment, the change in the transmission characteristic indicated by dotted lines in Fig. 3 can, for example, be caused by the simple act of someone opening a briefcase positioned on a desk. The raised lid of the briefcase results in a change in the characteristic. Similar extremely short term changes can be caused by the receiver 4 itself moving, or other objects moving such as doors opening, people moving, and the like. Clearly the transmitter 3 can also move. The foregoing establishes a very hostile environment within which the desired radio transmissions are to take place. In particular, there is no preferred channel or even a guaranteed channel within the 1 GHz band.

It would be possible to overcome the abovementioned difficulties by the use of highly directional antennae so as to eliminate all paths of transmission but the direct path. However, attempting to mechanically align such an antenna which was in turn affixed to a portable computer is commercially unattractive.

Fig. 4 illustrates in schematic form the general layout of a wireless LAN in accordance with a preferred embodiment of the present invention. A plurality of hubs 8 and mobile transceivers 9 are provided. The hubs 8 are interconnected by means of a backbone 10 which can take the form of either electrical conductors or optical fibre cable. As indicated by a dotted line in Fig. 4 the backbone 10 can constitute a loop. If desired, the backbone 10 can be connected to other computers 7 and, if desired, via a gateway 11 to the public

switched telephone network 12. In a typical arrangement, each office (or each office in each building of a campus) would be provided with a single hub 8 which would communicate with the, or each of the, mobile transceivers 9 in that room. Either the backbone 10 can extend over the entire area to be covered, or the area can be covered by the use of multiple gateways and multiple backbones. The effective range of the transceiver within the hub 8 is arranged to essentially cover only that room. The limited transmission range for the hub 8 creates a corresponding cell 13 as indicated by broken lines in Fig. 4. For a large room such as a lecture room in an educational institution, the length of the room can require that the room be provided with two hubs 8 in which case two partially overlapping cells 13 would be present within the one room.

As seen in Fig. 5, for the hub transceiver 8, a number of component blocks are provided. These take the form of a network interface 20, a buffer memory 21, a framing, forward error correction (FEC) and modulating section 22, a framing, forward error correction and demodulation section 23, an IF (intermediate frequency) system section 24, a mm-wave receiver 25, a mm-wave transmitter 26, and an antenna 27 which is sufficiently broad in its radiation pattern to illuminate the entire cell. The antenna 27 can achieve this result statically or dynamically (with electronic or mechanical beam steering). All these items are connected to, and are operable by, a control and timing section 28. In addition, all are powered by an AC mains operable power supply 29.

Equivalent portions of the mobile transceiver 9 are indicated by a designator having a magnitude higher by 10 in Fig. 5 and also in Figs. 6-9. The mobile transceiver 9 has a battery powered power supply 39. This is possible because of the use of low power gallium arsenide devices in the receiver 35 and transmitter 36, in particular.

It will be noted that the antenna 37 is preferably a steerable antenna which is electronically steerable by the control and timing section 38 so as to at least partially direct the transmissions to and from the mobile transceivers 9 towards the corresponding hub 8. A suitable antenna for this purpose is that disclosed in Applicant's Australian Patent Application No. ~~PL9739~~ 66100/94 entitled "A PLANAR ANTENNA" ~~(Attorney Reference 239045)~~, the contents of which are hereby

incorporated by cross-reference. This antenna improves the signal to noise ratio on the wireless link and attenuates delayed signals thereby improving multipath performance.

A more detailed block diagram of a portion of the transceiver 9 is illustrated in Figs. 6-9. In Fig. 6 the general arrangement of the transceiver 9 (excepting the terminal interface 30 and buffer memory 31) is illustrated. An intermediate stage of detail is given for the receiver 35 and transmitter 36, the receive intermediate frequency system 34 and receive demodulator 33 and the transmit intermediate frequency system 34 and transmit modulator 32. Full details of the modulation are given in Fig. 7 and of the demodulation in Fig. 8.

In Fig. 7, the transmit path framing, FEC and modulating section 32 of Figs. 5 and 6, is illustrated in detail. From the buffer memory 31 of Fig. 5 a binary data stream is applied to a CRC (cyclic redundancy check) Generate and Append block 40. The output of this block 40 or that of an End of Packet Pattern Generator 41 is selectively input to a rate 1/2 TCM (trellis coded modulation) Encoder 42. The output of encoder 42 is in turn input to a Di-bit Interleaver 43, the output of which is in turn input to a QPSK Encoder 44 which carries out differential encoding on a frame-by-frame basis. The output of QPSK Encoder 44 and a synchronising header generator 45 are combined in frame assembly and zero pad insertion block 46 so that the frames are assembled and four zero pads inserted so that six carriers are generated to each side of, but not coincident with, the centre frequency.

The assembled frames are then passed through an Inverse Fast Fourier Transform device 47 which uses a 16 point complex IFFT. The resultant signal is passed through Frame Serializer and Cyclic Extender block 48 to correctly sequence with 4 point cyclic extension the serial frames. The result is then passed via digital to analogue converters 49,50 to the intermediate frequency stage 34 of Figs. 5 and 6.

In the receive path in the Framing, FEC and Demodulating section 33 of Figs. 5 and 6, essentially the reverse procedures are carried out as illustrated in detail in Fig. 8. The received signal from the intermediate frequency stage 34 is passed through the analogue to digital converters 60,61 and thence to the cyclic extractor and frame assembler 62. The resultant signal is passed through the Fast Fourier

Transform device 63 to provide the essentially decoded signal. This signal is then simultaneously passed to a frame dis-assembler and zero pad remover 64 and to a synchronising calculator and detector 65 which provides start of message, end of message and symbol timing signals. These are passed to the control and timing unit 38 of Figs. 5 and 6.

The output of the frame dis-assembler and zero pad remover 64 is passed to a demodulator/detector 66 which carries out the necessary soft decision frame-by-frame differential demodulation and detection. The resulting output is passed to de-interleaver 67 and then to TCM decoder which is again a soft decision decoder. The decoder output is passed both to the buffer memory 31 of Fig. 5 and to CRC Accumulator and Checker 69. This latter device produces an error signal for the control and timing unit 38 of Figs. 5 and 6 if the demodulating/decoding has not correctly recovered the transmission data.

Turning now to Fig. 9, from the antenna 37 a schematically indicated bi-directional amplifier 71 leads via a filter 72 to an image rejection mixer 73. The preferred form of bi-directional amplifier 71 is that disclosed in Applicant's co-pending *International* ~~Australian~~ Patent Application No. PCT/AU94/00704 ~~PM-2445~~ entitled "A Bi-Directional Amplifier", the contents of which are hereby incorporated by cross-reference. Alternatively the bi-directional amplifier 71 can be realised by use of a separate transmit amplifier and a separate receive amplifier as illustrated connected between the antenna 37 and filter 72 by appropriate switches under the control of the control and timing unit 38 of Figs. 5 and 6.

The image rejection mixer 73 receives a 58GHz signal from a local oscillator (LO) unit 74. In the preferred form the first local oscillator (LO) is at a frequency of 58 GHz, resulting in an intermediate frequency band of 2-3 GHz. In the preferred embodiment illustrated in Fig. 9, this signal is obtained by doubling the output signal of a 29 GHz oscillator. It is also preferable to perform some form of frequency stabilization on this oscillator, either by using an external frequency discriminator as illustrated in Fig. 9, a stable internal resonator or some form of frequency/phase locked-loop.

The image rejection mixer 73 is connected to both the receive IF system 34 and the transmit IF system 34 and can be shared between them by the use of an appropriate switch again under the control of the

control and timing unit 38 of Figs. 5 and 6. The use of the filter 72 provides additional rejection of the image frequency.

From Figs. 6-9, it will be seen that the preferred form of modulation includes not only encoding but also fast fourier transforming, and its inverse. The transceiver 35,36 is preferably realised by means of one or more monolithic integrated circuits. Furthermore, in order to reduce power consumption in the mobile transceiver 9, the control and timing section 38 can power down each mobile transceiver 9 except when it is transmitting or receiving. This is determined by a polling scheme initiated by the hub transceivers 8. For example, the hub 8 can communicate with each mobile transceiver 9 in turn inquiring if any data is required to be transmitted or access to other parts of the LAN is required. This polling of the various stations can comprise one of a number of standard techniques such as time division multiple access, ALOHA or slotted ALOHA, timed token passing, grant request schemes or other applicable techniques.

The transmissions from the various transceivers 8 and 9 which comprise the network need not necessarily be at the same bit rate since some portions of the network need only a low speed of transmission (eg. printers) while others require a very high speed of transmission. This embodiment enables a variety of rates of transmission to be accommodated in a compatible network. This enables lower cost and/or low power consumption transceivers 9 to be used for printers or low data rate computing devices.

In order to provide a high speed bit transmission rate in the hostile radio environment as described above, at least two (and preferably three) techniques are used simultaneously. The first technique is to transmit over a relatively large number of parallel sub-channels within the available bandwidth so that each channel has a low bit rate but the total, or overall bit rate, is high. This spread, by increasing the symbol length, overcomes the problem of delay time and hence decreases the problems caused by inter-symbol interference.

The second technique involves the transmission of the data in small packets having some form of data reliability monitoring and/or enhancement such as Forward Error Correction (FEC). The length of the packet depends upon the method of data reliability enhancement and the hostility of the environment. Sufficiently small packets overcome the

problem of the rapid time change of the channel characteristics.

The third technique is interleaving (to be described hereafter) which is essentially a further data reliability enhancement. This technique improves the performance of many FEC schemes in overcoming the problems caused by nulls in the channel's frequency response.

In the most favourable environment, use of only ensemble modulation (the first technique) can be sufficient to produce an adequate result. However, such environments are rarely encountered and therefore, in practice, the second technique should be employed in combination with the first technique.

The initial form of the second technique is data reliability enhancement by automatic repeat request (ARQ). The maximum permissible packet length able to be chosen is that which will ensure a practical probability of error free transmission. As the hostility of the environment increases, either channel sounding or a redundancy arrangement such as forward error correction (FEC), and/or data redundancy, and/or permutation modulation should be also used. If necessary, both channel sounding and redundancy technique(s) can be used.

In relation to the first of these techniques, typical time delays due to multipath transmission are of the order of 50 ns because of the dimensions of typical rooms. At a desired bit rate of the order of 100 Mbit/s, this indicates that the bit period is 10 ns which is only 20% of the delay time. However, if the transmission is divided into, say, twelve sub-channels, then in order to achieve a bit rate of 100 Mbit/s overall, this implies that each channel must have a bit rate of approximately 8.3 Mbit/s. If 12 bits are encoded and sent as a symbol, then the symbol time is of the order of 120 ns which is greater than the delay time. The choice of the optimium number of sub-channels depends on the enviornment.

In relation to the second technique, because of the fading channel, not all the sub-channels can be expected to transmit successfully. For this reason data error correction is provided. This takes a number of forms. The first is redundancy sufficient for the detection of errors so that there may be subsequent re-transmission of at least selected data in which those passages of information not correctly received are re-sent. The re-transmission is not necessarily

over the same sub-channel or channel. The second is forward error correction which has a redundancy sufficient for non-interactive correction. A third is permutation modulation such as multi-tone amplitude shift keying which has built-in redundancy. Any of these techniques allow the demodulator to correct for a relatively small percentage of errors in the received bits.

The preferred type of modulation in each sub-channel is multi-level modulation of carrier amplitude and/or phase (mQAM). The modulation family mQAM includes amplitude shift keying (ASK), multi-level ASK (mASK), permutation modulation, binary phase shift keying (BPSK), multi-level phase shift keying (mPSK), amplitude phase keying (APK), multi-level APK (mAPK) and the like.

Transceivers 9 for devices such as printers which require a lower bit rate transmission can use the techniques which give a lower spectral efficiency such as amplitude shift keying (ASK).

A variant of ASK over an ensemble of carriers is called permutation modulation. In this scheme, the transmission is m-ary where a transmitted symbol can encode $\log_2 m$ binary digits. There is an alphabet of m symbols allocated to the channel. Each symbol transmitted has a built in redundancy so that if several of the symbols are received in error due to the poor nature of the corresponding part of the channel, a correct decision can still be made as to which of the allowed symbols was transmitted.

A choice of the symbols with the appropriate orthogonality can be made using a number of well known information theory techniques or by a computer search for the appropriate codes. Due to the high redundancy and limited bandwidth efficiency of permutation modulation, this system does not yield a high spectral efficiency (expressed as bits/Hz). For the system of the illustrated embodiment this can be lower than 0.25 bit/Hz. It is, however, relatively simple to implement and so is desirably used in a lower cost, lower bit rate transceiver 9 for printers, for example, which are compatible with the higher performance embodiments described below.

Another embodiment of the multi-carrier scheme is to phase modulate each carrier using a phase shift keying (PSK). In simple embodiments this is binary phase shift keying (BPSK) where two phase options are transmitted or quadrature phase shift keying (QPSK) where

13

a *Phase*

four options are transmitted. Any higher number can be transmitted as required.

In the BPSK embodiment, incorporating forward error correction, the incoming binary data stream at a bit rate "b" is encoded using a conventional forward error correction scheme such as, but not restricted to Reed-Solomon or convolutional coding. Such coding increases the number of bits to be transmitted by a factor "r" which is the reciprocal of the code rate. The encoded stream, at a bit rate b.r, is then split into "p" parallel paths, and each path used to BPSK modulate a separate carrier in the ensemble, giving an effective symbol duration, on the radio link, of p/(b.r.) seconds.

The resulting signal is then transmitted over the channel and received by the other unit with some sub-channels error free and others with a potentially substantial error rate, due to the frequency selective nature of the channel.

The received carriers are demodulated and the individual bit streams combined (or aggregated) to form an encoded data stream with possible errors (mainly from the bad sub-channels) which is then decoded by a device (such as a Reed-Solomon or Viterbi decoder). Any errors in the received signal are normally completely corrected by this decoding process.

Additionally, a weighting can be given as to the confidence of the accuracy of the output of each BPSK demodulator based upon the amplitude of the received carrier. This weighting can be used as an additional input to the decoding device to determine which bits are more likely to be in error and to increase the performance of this device in correcting as many errors as possible in the transmission.

It is also possible to use combined coding and modulation schemes such as trellis-coded-modulation (TCM) to give improved bandwidth efficiency and improved error correction capability.

It is also possible to use multiple level phase shift keying modulation on each of the carriers transmitted and a corresponding demodulator on the receiver. This will give improved bandwidth efficiency and therefore allow much higher data rates to be transmitted through the channel for the same compatible bandwidth. This option allows higher bit rate units to occupy the same spectrum as the lower bit rate transceivers but in a compatible manner. The increased

spectral efficiency is acquired at the cost of increased complexity in the modulators and demodulators, together with some degradation of error performance.

As referred to above, link data interleaving schemes can be used, in this system, to further improve the error correcting performance of FEC codes which distribute the contribution of individual data elements over fewer carriers than the total number in the ensemble. Link data interleaving schemes do this by distributing the encoded data between the carriers in such a way that the correlation in error probability of those carriers associated with any given element of uncoded input data is minimised. On average, this corresponds to maximising the minimum frequency spacing between those carriers.

For example, with a 5-bit constraint length, half rate trellis coded QPSK modulation on the carriers of a 12 carrier ensemble, a suitable interleaving scheme is -

Carrier number (1-12) modulated by successive encoder output di-bits:

1, 3, 5, 7, 9, 11, 2, 4, 6, 8, 10, 12, 1, 3, ... etc.

Such an interleaving scheme is typically implemented by means of demultiplexers, shift registers and multiplexers in substantially conventional fashion.

The above will improve the error rate performance of the system, however, it will not eliminate all errors in all cases. To overcome any residual errors in the system an additional error correction layer, such as cyclic redundancy checked (CRC) automatic repeat request (ARQ),can be used. This error correction layer requests the re-transmission of those symbols which are believed to be in error. This re-transmission can occur over the same frequency channel, or a request can be made to the control and timing section to shift the entire frequency channel by some predetermined amount, or to change antenna characteristics such as polarisation, to increase the probability of error free transmission.

Because of the highly time variable nature of the transmission channel, the transmitted data is divided into packets of short duration (typically 100 microseconds). During this short time period it is satisfactory to assume that the transmission characteristics are essentially stationary. Before transmission of a packet of data, it is

possible to use a channel selection technique to reduce error rates. One channel selection technique is to channel sound prior to transmission of the packet. If necessary, this allows the data rate to be reduced if a particular sub-channel or channel is found to be degraded.

As illustrated in Figs. 7 and 8 the preferred method of generating and demodulating multi-carrier modulation schemes uses a device capable of performing Fast Fourier Transforms (FFTs) and Inverse Fast Fourier Transforms (IFFTs) on complex data at high speeds. Such a device is described in Australian Patent No. 610,934 entitled "A Transform Processing Circuit" granted to the present applicants, the disclosure of which is hereby incorporated by cross-reference. In the example shown in Figs. 7 and 8, 16 point fast Fourier transformation is used.

Improved performance can be obtained by using cyclic extension by means of circuit 48 and cyclic extraction by means of circuit 62 in conjunction with the fast Fourier transformation. Cyclic extension is a technique for enhancing the multi-path tolerance of FFT-based ensemble modulation schemes by reducing the degradation of sub-channel orthogonality produced by channel delay spread effects and demodulator timing errors. It consists, at the modulator, of extending the time duration of individual multicarrier symbols by appending, to the FFT output frame, a copy of that frame, then truncating the combination to the desired length. The length of the extension is a compromise between tolerance to multipath induced intra-symbol interference and the reduction of channel spectral efficiency. It preferably corresponds to the time interval over which the channel impulse response has substantial energy.

At the extractor and frame assembler 62, an essentially uncorrupted multicarrier symbol is excised (by cyclic extraction) from the potentially distorted incoming extended symbol, whose ends may be corrupted by the extended impulse response of a multipath channel. This excised symbol is then used in the FFT based demodulation process. For example, when using a 16 point FFT, a cyclic extension length of 4 points can be used.

These processes can be effectively implemented by a slight extension of the frame assembly/disassembly mechanism required for the

FFT interface. A related (but more computationally intensive) process is that of "tapering" or "windowing", whereby the amplitude of the multicarrier symbol is varied over part of the symbol time in order to reduce the mutual cross-talk of sub-channels more than a few carrier spacings apart in frequency.

When using multi-carrier schemes it is not always desirable to occupy the full band and some carriers need not be transmitted. For example, when using an FFT device 63 the analogue (reconstruction/ anti-aliasing) filter selectivity requirements, for given adjacent channel suppression/rejection, can be relaxed by using a larger transform whose higher frequency bins are zero-filled in the modulator and ignored in the demodulator. This corresponds to not generating the higher frequency (out-of-band) carriers at the transmitter and ignoring any received energy at those frequencies, so the FFT provides (subject to dynamic range considerations) a significant part of the band edge selectivity. Zero insertion can also be used to remove the band-centre carrier (DC at baseband) to reduce susceptibility to DC offset drifts in the system. For example, when using a 16 point FFT device 63 only 12 carriers are preferably used.

As shown in Fig. 6, a device 65 is required to synchronize the receiver to the incoming data. This device can, for example, compare this incoming data to the receiver's timing signals, calculate the difference in sumbol and bit times and pass this information to the control and timing unit which would then perform the appropriate corrections to achieve synchronization or zero difference.

A preferred synchronisation scheme, having multipath tolerance commensurate with ensemble modulation and sharing of the FFT hardware, determines multicarrier symbol timing and gross local oscillator frequency difference by measurement of the relative phases of several carriers present in the message header generated by generator 45 and comparision of these with the known phase relationship of the transmitted carriers at the beginning of the header transmission.

The IF systems 34 are shown in Fig. 6 and consists of an I,Q up converter for the transmitter and I,Q down converter for the receiver. The second LO units of the IF systems 34 are tuned over the band 2-3 GHz and this allows the conversion of the signals to and from baseband. In some embodiments it is preferable to provide tuning of

the carrier frequency by varying the frequency of the first local oscillator 74 (Fig. 9) and in others by varying the frequency of the second local oscillator in the IF system. It is possible to share some of the components in the transmit and receive IF systems.

The foregoing describes only some embodiments of the present invention and modifications can be made thereto without departing from the scope of the present invention. For example, interleaving and bit reversal of the transmitted data to decrease the received error rate can be accomplished by utilizing the bit reversal inherent in the FFT conversion. Also the antenna 37 can utilize polarisation diversity to improve reception.

One arrangement for the simultaneous operation of low bit rate transceivers and high bit rate transceivers is to allocate, say, half the available (high bit) channel to the low bit rate transceivers. Thus, the low bit rate transceivers utilize only half of the available bandwidth and a hub can transmit data at the low rate to two low bit rate transceivers at the same time. Thus the same hub hardware is used for both high bit and low bit rate transmissions.

It will be clear to those skilled in the art that the LAN need not incorporate hubs 8 since the mobile transceivers 9 can transmit to, and from, each other directly within the predetermined cell range. Such a LAN is termed a peer-to-peer LAN.

Similarly, the hubs 8 although described as being interconnected by electric cable and/or optical fibre, can also be inter-connected by a radio or infra-red link. The link can form a part of the backbone 10 or constitute the inter-hub communication link.

18